BENCH MARK

April 2016 issue . . .

THE INTERNATIONAL MAGAZINE FOR ENGINEERING DESIGNERS & ANALYSTS FROM **NAFEMS**

The HPC Edition

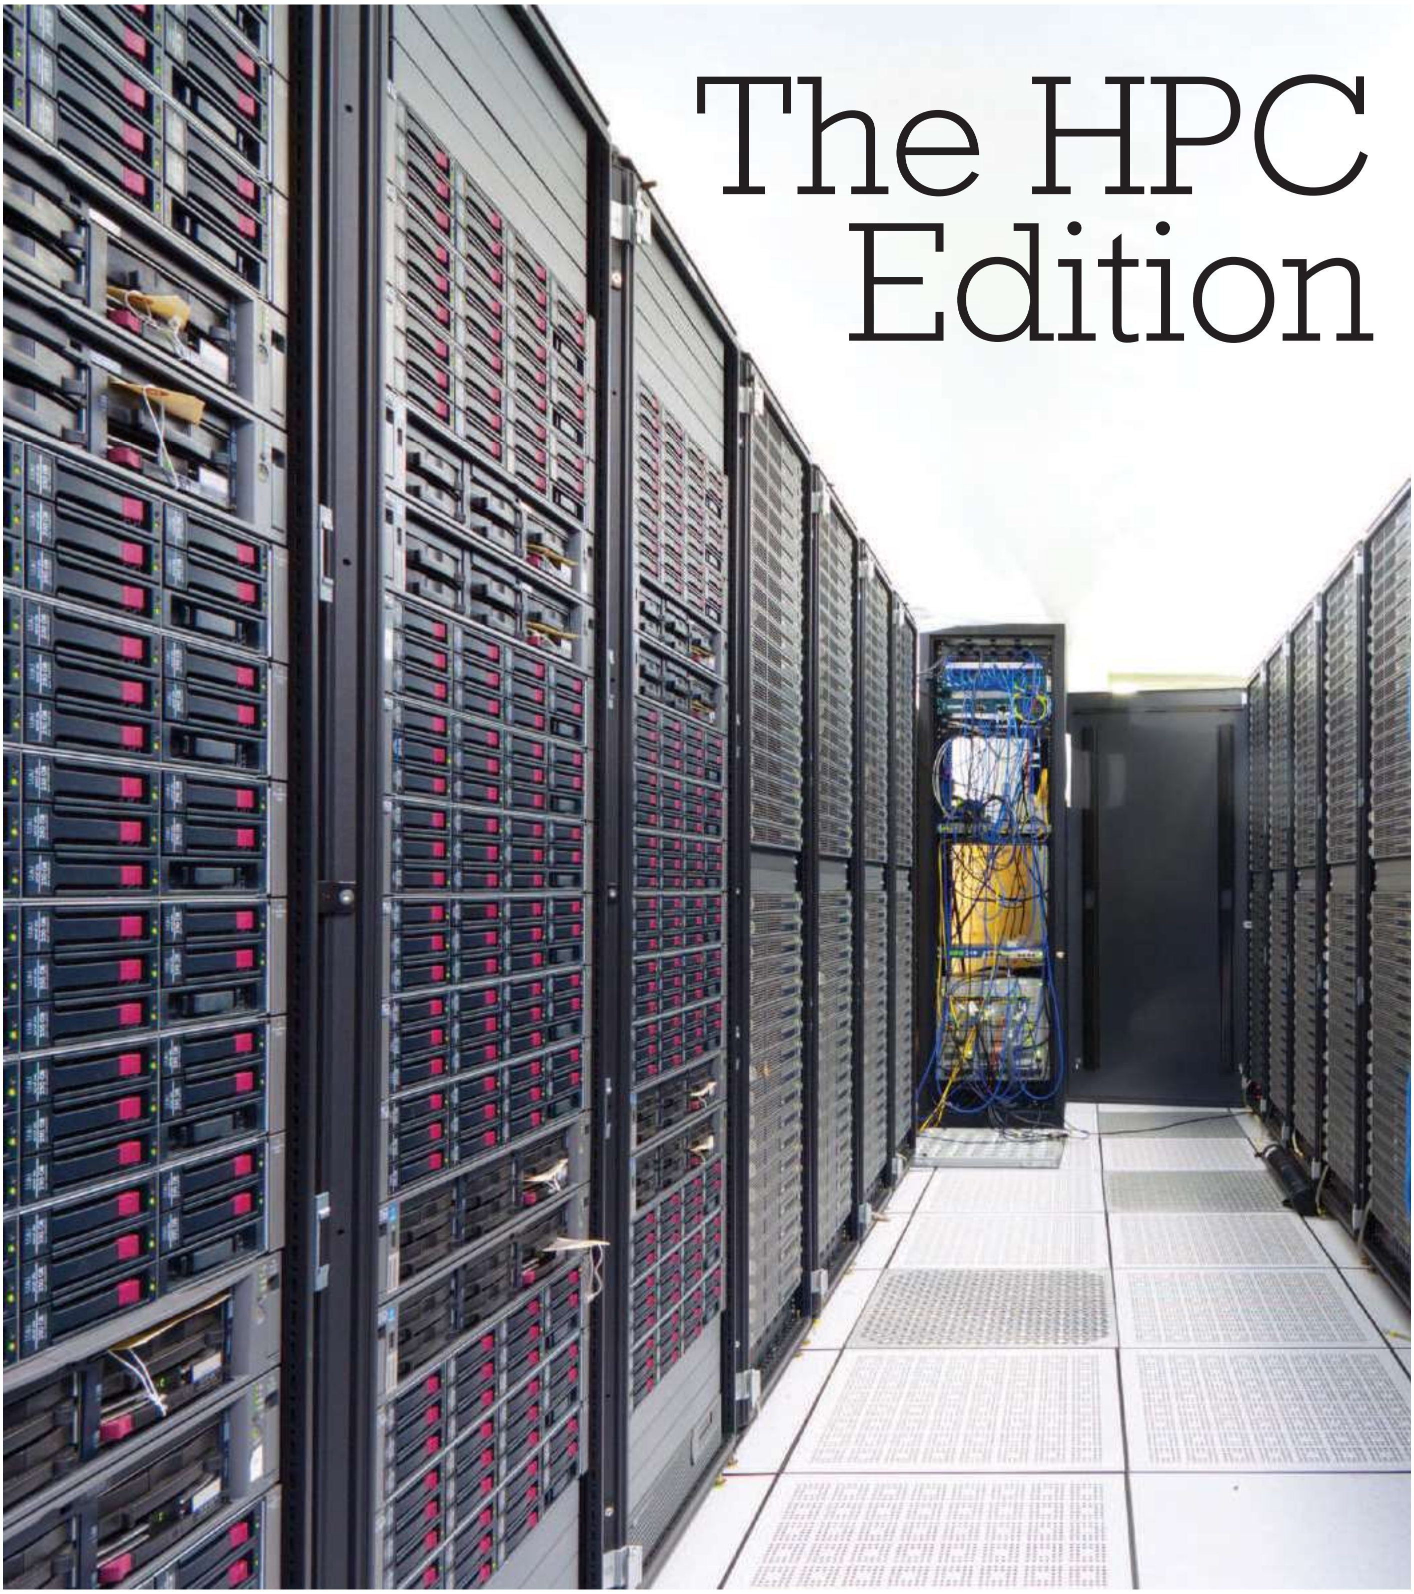

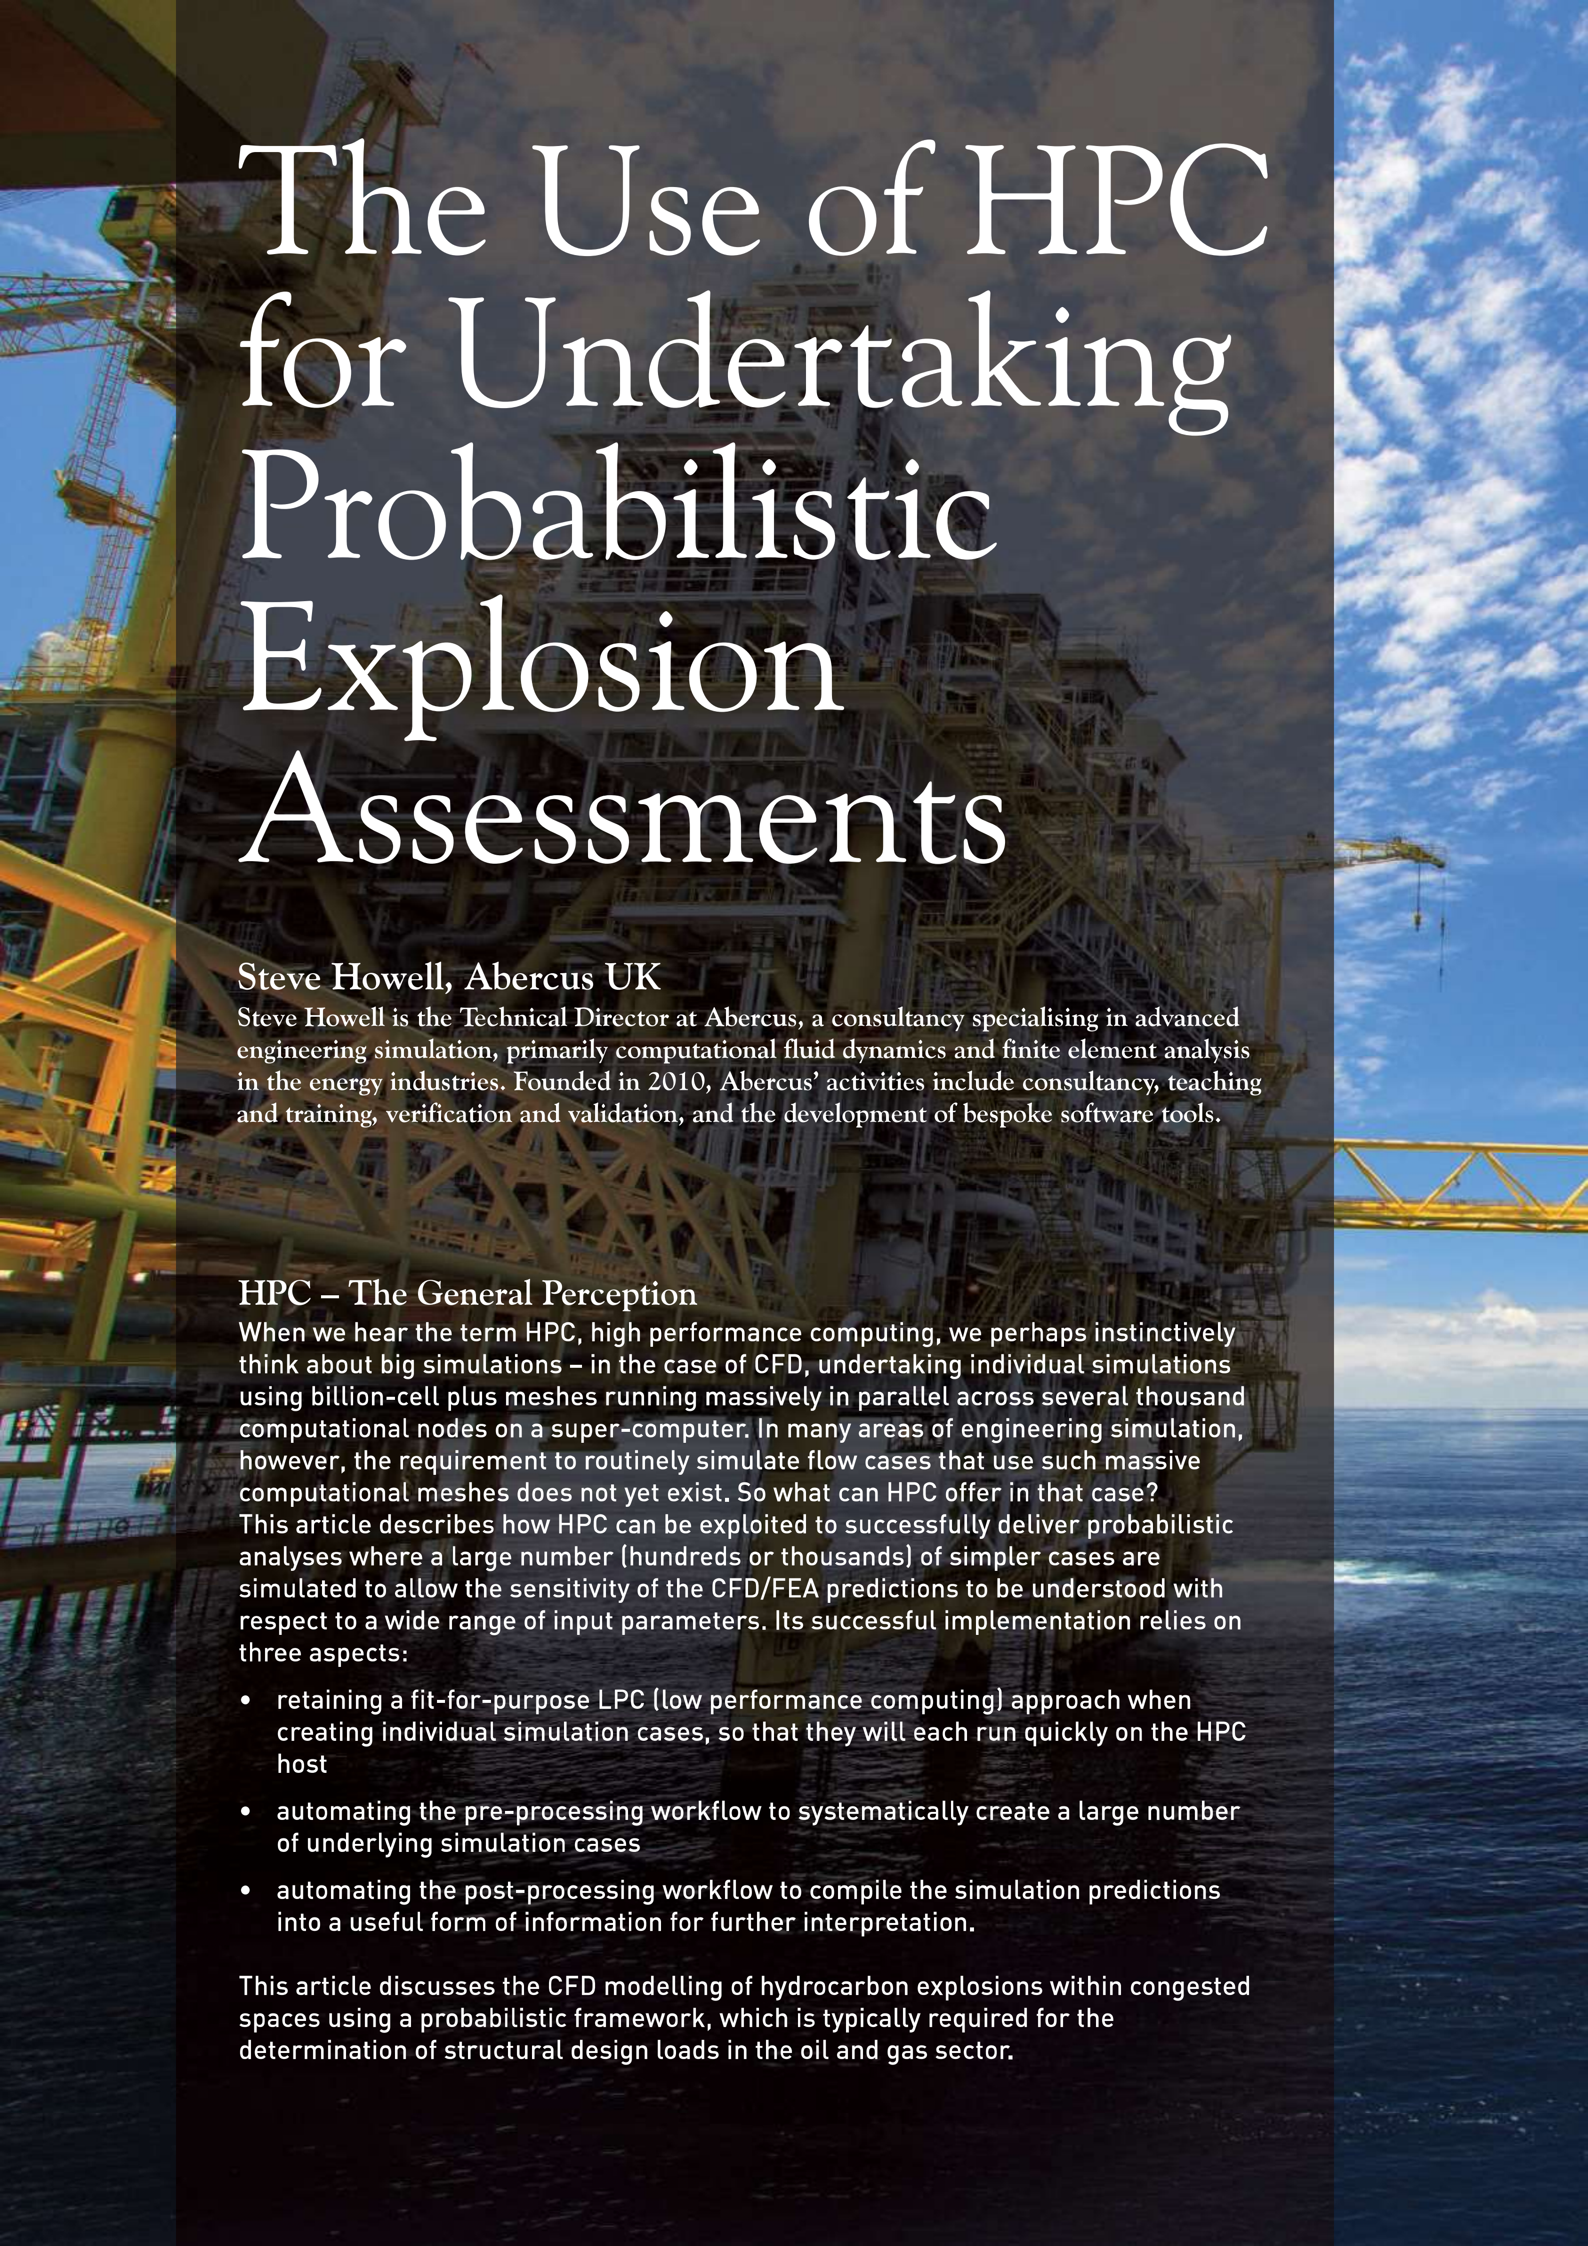

The Use of HPC for Undertaking Probabilistic Explosion Assessments

Steve Howell, Abercus UK

Steve Howell is the Technical Director at Abercus, a consultancy specialising in advanced engineering simulation, primarily computational fluid dynamics and finite element analysis in the energy industries. Founded in 2010, Abercus' activities include consultancy, teaching and training, verification and validation, and the development of bespoke software tools.

HPC – The General Perception

When we hear the term HPC, high performance computing, we perhaps instinctively think about big simulations – in the case of CFD, undertaking individual simulations using billion-cell plus meshes running massively in parallel across several thousand computational nodes on a super-computer. In many areas of engineering simulation, however, the requirement to routinely simulate flow cases that use such massive computational meshes does not yet exist. So what can HPC offer in that case?

This article describes how HPC can be exploited to successfully deliver probabilistic analyses where a large number (hundreds or thousands) of simpler cases are simulated to allow the sensitivity of the CFD/FEA predictions to be understood with respect to a wide range of input parameters. Its successful implementation relies on three aspects:

- retaining a fit-for-purpose LPC (low performance computing) approach when creating individual simulation cases, so that they will each run quickly on the HPC host
- automating the pre-processing workflow to systematically create a large number of underlying simulation cases
- automating the post-processing workflow to compile the simulation predictions into a useful form of information for further interpretation.

This article discusses the CFD modelling of hydrocarbon explosions within congested spaces using a probabilistic framework, which is typically required for the determination of structural design loads in the oil and gas sector.

Modelling Explosions within Congested Spaces

Whilst hydrocarbon explosions can, in principle, be simulated using general-purpose CFD codes, this is not common practice (at least in the oil and gas sector). Instead, a number of explosion-specific CFD codes have been developed, including FLACS by Gexcon, EXSIM by ComputIT, and AUTOREAGAS by Ansys. Each of these codes has been developed to simulate deflagrations (subsonic explosions) and they all follow a similar fit-for-purpose LPC approach.

The distributed porosity approach for capturing the effects of small scale congestion

Within congested spaces, the amount and spatial distribution of small-scale obstructions such as structural steelwork, equipment, pipework and pipe-supports, and cable trays, which are abundant across the topsides of any offshore platform, affect the background ventilation flow and the dispersion behaviour following a release. The level of congestion also has a major impact upon the intensity of the explosion, due to the Shchelkin mechanism whereby the turbulence generated by the congestion causes an increase in the rate of combustion at the flame front [1, 2].

In principle, it is possible to capture the small-scale congestion explicitly within a CFD mesh and undertake a transient combustion simulation to model explosion events using a general-purpose CFD code and capture the Shchelkin mechanism from first principles. This approach, however, is typically prohibitive, both in terms of the cell count of the CFD mesh and the small time-step that would be required to satisfy CFL constraints. Whilst an HPC computing resource may allow this first-principles approach to be pursued for, perhaps, a handful of explosion cases within a reasonable time frame, this is unlikely to be appropriate for a probabilistic approach in the near future, where hundreds or thousands of individual explosion cases may need to be simulated.

In contrast, the explosion-specific CFD codes are based upon a distributed porosity approach, similar to the ACE method where a relatively coarse CFD mesh is used and the effect of the sub-grid congestion is captured by allocating equivalent resistance and generation source terms in the momentum and turbulence equations respectively [3]. Within the combustion model, semi-empirical correlations are used to predict the effect of the Shchelkin mechanism upon the explosion behaviour.

Simplification of the near-jet region for a high pressure compressible release

For high-pressure compressible releases, the near-field region of the emerging jet has high momentum and is not significantly affected by the background ventilation flow. Instead of explicitly capturing this within the CFD model, it is usually advantageous to separate the near-field and use another method to predict the flow in this region, a classical calculation or a separate axisymmetric CFD approach. This approach will typically not significantly affect the predicted dispersion behaviour in the far-field [4] and is adopted by each of the explosion-specific CFD codes.

Structured orthogonal mesh

Each of the explosion-specific codes is based upon the structured orthogonal mesh approach. Whilst this can allow the governing equations to be solved more efficiently than with general-purpose CFD codes which necessarily are programmed to solve on unstructured meshes, this does mean that the actual CFD geometry will resemble a *Lego* model, as shown in Figure 1.

Designing for Explosion and the Probabilistic Approach

The sequence of events leading up to an explosion

When modelling a hydrocarbon explosion event it is necessary to understand the deterministic sequence of events that lead up to the explosion, including the background ventilation flow behaviour due to the prevailing wind condition and any HVAC (Heating Ventilation and Cooling) systems that may influence the flow, the discharge and dispersion behaviour following a loss of containment of flammable material, and the explosion dynamics following the subsequent ignition of the accumulated cloud of flammable material, as shown in Figure 2.

The probabilistic approach and the compilation of exceedence data

When designing for explosion it is necessary to determine a suitable basis for the design explosion loads that a structure must withstand. A worst-case approach considering just a few explosion scenarios following large release events is typically overly conservative, and the corresponding explosion loads may be well in excess of what can be realistically designed for. Instead it is usual to adopt a probabilistic approach where explosion scenarios following a large number of releases representing the whole spectrum that could occur (from small through to large/full-bore releases) are simulated for a range of background ventilation conditions.

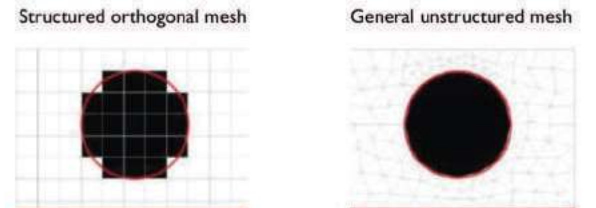

Figure 1: Limitation of the structured orthogonal meshing approach for capturing non-orthogonal geometries

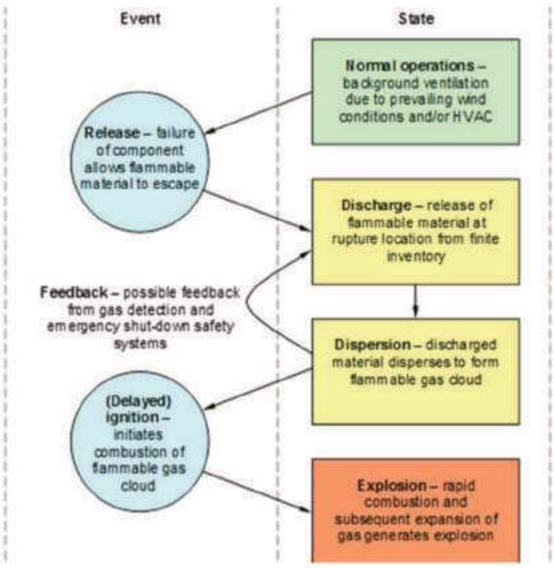

Figure 2: The deterministic sequence of events leading up to an explosion

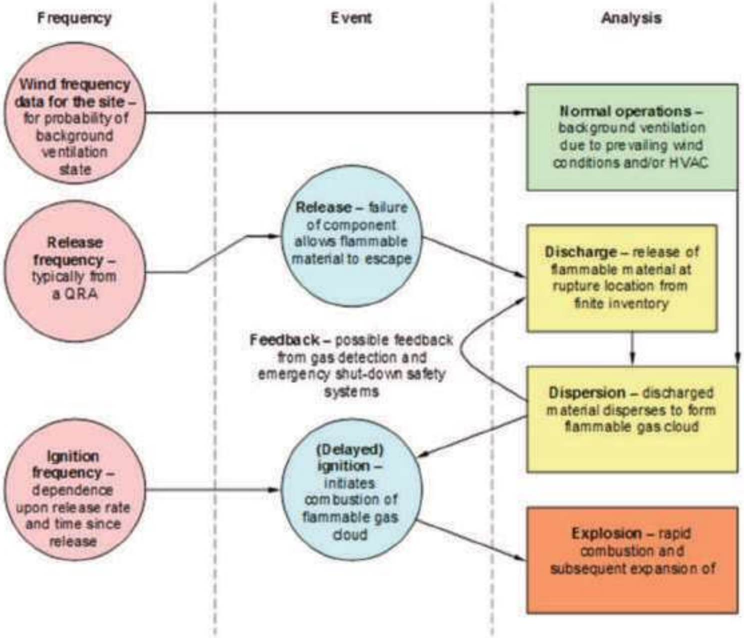

Figure 3: The probabilistic framework

The probabilistic methodology used in the oil and gas sector is described in NORSOK Standard Z-013 [5]. Critically, the frequency of occurrence for each simulated scenario is calculated at each stage of the sequence, using wind frequency data for the background ventilation scenarios, together with the release frequency for the discharge/dispersion scenarios, and the ignition frequency for the explosion scenarios, as shown in Figure 3.

Upon completion of the simulated dataset, the predictions are compiled to form exceedence curves which describe the probability that an explosion event of a particular magnitude will occur. Using the exceedence curves, the explosion load corresponding to the risk acceptance/frequency cut-off criteria is determined, as shown for the example in Figure 4 – here the acceptable level of risk is taken to be a 1 in 10 000 year explosion event, which corresponds to a frequency cut-off of 10^{-4}/yr.

The challenge – keeping the number of simulated cases in the dataset manageable

The challenge to date has been to keep the dataset of simulated scenarios manageable, so that the simulations can be completed within a reasonable timescale, typically on high-specification desktop workstations. The issue is that for any single explosion sequence there are a wide range of parameters that will influence the behaviour at each stage in the sequence, as shown in Figure 5. Consequently, for a typical facility there could be hundreds of thousands of possible dispersion scenarios and several million possible explosion scenarios to consider within the probabilistic framework. To simulate them all is simply not feasible. On a modest resource of, say, 100 compute nodes this would take several years to complete. Even with an HPC resource of, say, 10 000 compute nodes this could take several months.

Simplifying assumptions

It is necessary to make assumptions and recognise similarities and symmetries to identify a reduced number of key representative dispersion and explosion scenarios to simulate. All of the possible scenarios (and their probability of occurrence) are then allocated against one of the scenarios actually simulated. Provided any assumptions introduced are conservative, the probabilistic analysis should also remain conservative.

Ventilation and discharge/dispersion simplifications

Dispersion scenarios may not be explicitly simulated for all wind conditions. Instead, it is common to assume that low wind speed conditions prevail, which should be conservative since higher wind speeds tend to provide better ventilation and disperse the released hydrocarbons more quickly. If this approach is found to be too conservative, higher wind speed conditions can be simulated too.

Dispersion scenarios may not explicitly be simulated using CFD for all of the release rates required by the NORSOK Standard [5] (which requires that a minimum of nine different hole sizes/release rates are considered for each release). Instead, simulations may be undertaken for selected release rates across the range required by the standard, with the dispersion behaviour of intermediate releases determined through conservative interpolation from the simulated cases.

Dispersion scenarios may not be explicitly simulated for all possible release directions. Instead, symmetries are exploited where the local geometry makes it possible to assume that a release in one direction is also representative of releases in other directions – particularly for smaller releases where the flammable envelope may not interact significantly with the

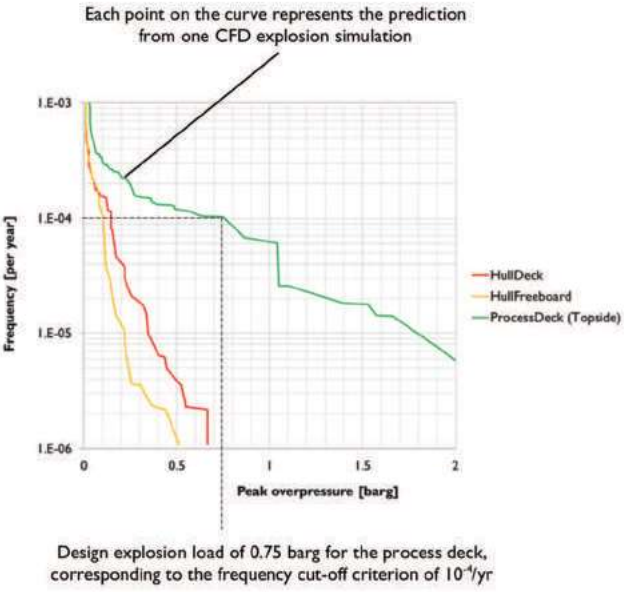
Each point on the curve represents the prediction from one CFD explosion simulation

Design explosion load of 0.75 barg for the process deck, corresponding to the frequency cut-off criterion of 10^{-4}/yr

Figure 4: Example exceedence curves for explosion loads from a recent probabilistic study, used to determine the structural design loads

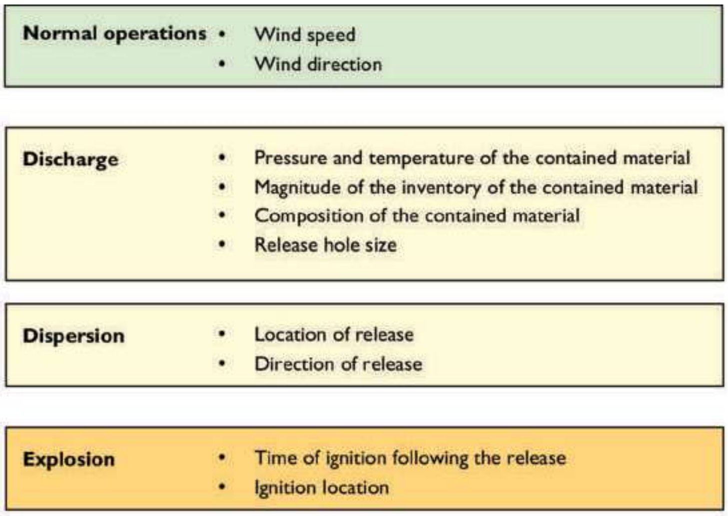

Normal operations	• Wind speed • Wind direction
Discharge	• Pressure and temperature of the contained material • Magnitude of the inventory of the contained material • Composition of the contained material • Release hole size
Dispersion	• Location of release • Direction of release
Explosion	• Time of ignition following the release • Ignition location

Figure 5: Parameters that influence each stage in the explosion sequence

surrounding geometry. The simulated release directions are usually restricted to the principal orthogonal directions, to minimise any numerical diffusion issues associated with the structured orthogonal mesh approach (although it should be noted that at least one of the explosion-specific code providers is developing an adaptive mesh capability which could relax this restriction).

Together, these simplifications allow the total number of transient dispersion cases that need to be actually simulated within a probabilistic study to be significantly reduced, typically in the range of a few hundred to a couple of thousand simulations.

Explosion scenarios

For each transient dispersion simulation, the hydrocarbon concentration will vary spatially throughout the flammable envelope and will also vary with time. To account for this, the EI ignition correlations [6] define a time-based ignition probability for six time-windows covering the duration of the release. If a snapshot of the dispersing cloud is ignited for each time-window in three separate locations (which is typically a minimum requirement) this corresponds to 18 individual explosion simulations for each dispersion simulation. Even with the reduced number of simulated dispersion cases, maintaining a direct coupling between the dispersion and explosion stages and modelling each of these possibilities outright would lead to tens of thousands of individual explosion simulations. To reduce this computational overhead, the NORSOK Standard allows the dispersion and explosion stages to be decoupled. Specifically, each of the spatially-varying flammable clouds from the dispersion stage can be represented by a homogeneous cubic cloud with an equivalent stoichiometric volume at the explosion stage.

This simplification, which was necessitated by the limited computing power available in the 1990s and early 2000s, has been criticised. The envelope of the equivalent stoichiometric cloud, by definition, will be physically smaller than that for the original spatially-varying cloud, so the latter will typically envelop more congestion than the equivalent stoichiometric cloud. Since it is known that congestion impacts upon the severity of an explosion, the two approaches are not necessarily consistent. Despite these concerns, the method is allowed by the NORSOK Standard and is widely used in practice [7]. It allows the number of explosion simulations to be significantly reduced, typically to just a couple of hundred.

Using HPC Computing Facilities for Probabilistic Explosion Studies

By exploiting the simplifying assumptions described above and minimising the number of simulated cases within the underlying dataset, the dispersion scope may take around one month to complete and the explosion scope may take a few days to complete on a modest resource comprising, say, 100 compute nodes. The entirety of a typical study may take a couple of months, allowing for the initial CFD model build, meetings and reporting.

Clearly there is an opportunity to exploit HPC in order to reduce the time required to complete probabilistic assessments. This opportunity has been identified by Gexcon, the providers of FLACS, who have recently launched a new HPC service for explosion modelling in partnership with EPCC, one of Europe's leading supercomputing centres based at the University of Edinburgh [8, 9]. The partnership has been trialled since 2014 and initial feedback has been overwhelmingly positive.

Simulation data management (SDM) and the automation of the simulation workflow

With such a large number of individual simulations within a probabilistic study, the success of the HPC opportunity is greatly enhanced by a robust simulation data management (SDM) framework with the ability to automate the simulation workflow, in terms of setting up the individual simulations that need to be undertaken, running them automatically in batch mode, and then compiling the CFD predictions automatically into the exceedence data from which the structural design loads can be determined.

Abercus has recognised this and has developed the EXCGEN software to automate this workflow. Gexcon is currently developing a tool called RISK and other practitioners have undoubtedly developed their own in-house codes which incorporate similar functionality.

Whilst a major benefit of these SDM tools is to automate the workflow to capture the HPC opportunity, there are additional benefits that should be recognised [10]:

- These tools provide a robust, consistent method for the implementation of the NORSOK Standard Z-013. Currently there is scope for differing interpretation of the probabilistic approach outlined in this standard, so it may be that different approaches are adopted by different practitioners. SDM tools like EXCGEN and RISK could allow a common approach to become more easily adopted, provided the underlying implementation is openly documented.
- These tools potentially allow for sharing and democratisation of analysis data. EXCGEN's underlying file system is generated from, but separate to, the underlying CFD files, which allows the information to be easily shared. Using EXCGEN, the sensitivity of the exceedence data to many of the probabilistic assumptions can be investigated on-the-fly, in the company of the wider design team. Without an SDM tool, these sensitivities cannot be explored so easily.
- The automatic compilation of the exceedence data can be extended to allow 3D risk assessment where, for example, the spatial variation of an explosion load can be presented across a structural target of interest, rather than just a single worst-case load that is read from a traditional exceedence curve. Traditionally only exceedence curves have been provided to the structural engineering team, at least partly because there has not been a tool that has been capable of compiling the vast amount of data generated during a probabilistic assessment in a useful and easily transferable format.
- The SDM tools can be extended to allow the outputs from an explosion CFD case (or a pseudo-case derived from the compiled exceedence data) to be automatically mapped on to an FEA model so that the associated structural response can be simulated.

Additional value by exploiting HPC computing with the current probabilistic approach

By using HPC for probabilistic explosion assessments, the typical project timeframe of around two months can be shortened, perhaps halved. The oil and gas industry can be a slow beast, however, with projects often lasting for months or even years, so the benefit of saving perhaps one month from a probabilistic assessment could easily be missed in the scale of the rest of the overall project. Instead of simply shortening the time taken to complete an assessment, there may be better ways to exploit HPC which could add more value.

Optimisation of the layout of a facility

At present, because of the lengthy simulation times, the dispersion and explosion scenarios are often simulated for just one instance of the geometry of the installation. The entire probabilistic study is, therefore, an assessment of that single geometry instance. If the simulation time can be shortened to just a few days for the dispersion/explosion stages then several geometries could be considered within the typical two-month project window. This could provide significant value to the project team and allow some optimisation of the facility layout,

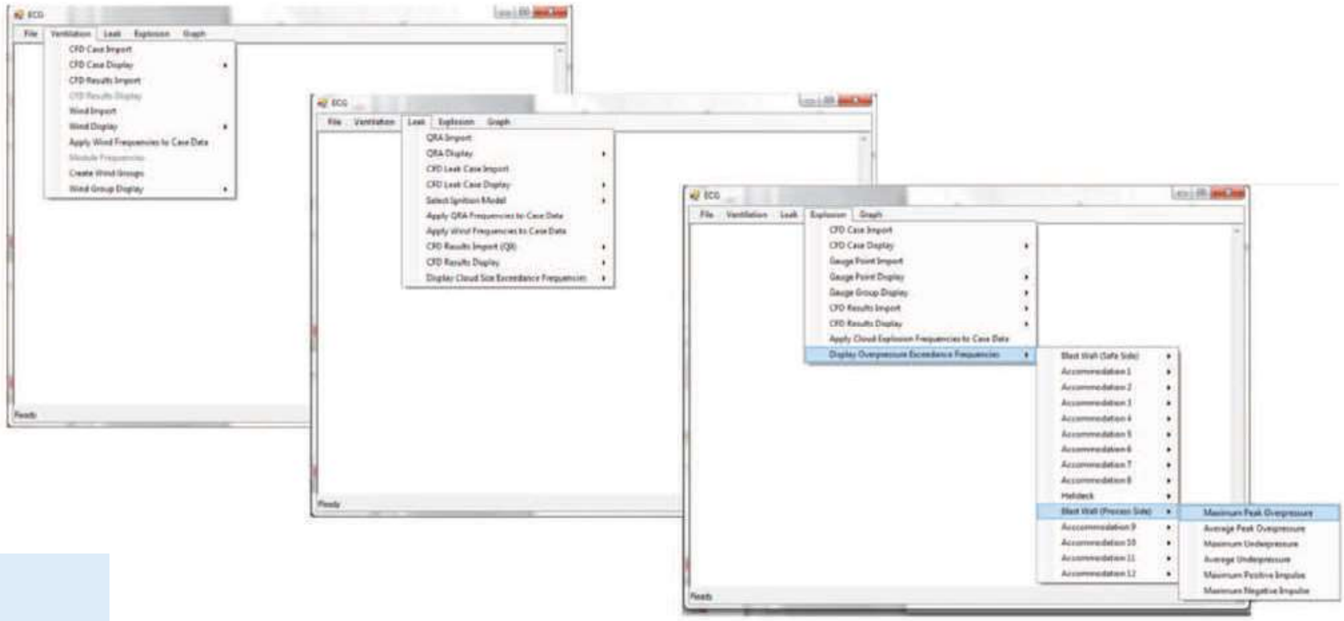

Figure 6: The EXCGEN GUI, used to automate the simulation workflow for probabilistic explosion assessments

for example, with respect to the spacing between decks, the positioning of blast walls and their optimum dimensions, the optimal layout of/spacing between process modules within a given footprint, and an assessment of whether decks should be grated or plated, or some proportion of both.

Probabilistic structural response

Traditionally the structural design loads are determined entirely separately to the structural design team that actually makes use of them. Despite the huge amount of information generated as part of a probabilistic explosion assessment, the interface between the two teams has traditionally been limited to a few exceedence curves which provide the worst-case load on each structural target. This interface may be excessively conservative on the basis that the load from the exceedence curve is assumed to act uniformly across the structural target.

By capturing a more realistic spatial distribution of the explosion loading across the structure this conservatism may be reduced. The identification of a representative *dimensioning* explosion event from the dataset underlying the probabilistic study, somehow corresponding to the risk acceptance/frequency cut-off criteria, is one possible approach but the selection of this event is not a trivial task [10]. The construction of a single pseudo-event corresponding to the frequency cut-off of interest, automatically constructed from the underlying dataset, is another promising approach but requires further investigation [10]. Both of these methods rely on the identification of a single representative explosion event that is imposed upon an FEA model to investigate the corresponding response of the target structure and determine the damage/deflection associated with the risk acceptance criteria used to construct the explosion event.

With HPC, together with the development of the SDM tools, an alternative systematic one-to-one coupling between the explosion CFD code and a structural FEA code is possible, where the predicted loads for every simulated explosion scenario are automatically mapped on to an FEA model for each structural target of interest and the structural response simulated directly. The big advantage of this approach is that it is not necessary to make any assumptions about the representative explosion event. As a consequence of the direct one-to-one coupling, the probability of occurrence of each simulated structural case is equal to that of the corresponding explosion case, so it becomes straight-forward to automatically construct contour plots for the deflection corresponding to the risk acceptance criteria/frequency cut-off of interest [10].

This probabilistic structural response approach is currently being investigated by several parties [10, 11]. On a recent project Abercus found that the additional computational overhead to simulate each of the FEA structural response cases was fairly modest compared with the timescale for the rest of the project and, in particular, the simulation time required for the dispersion stage. With an HPC resource, together with a tool like EXCGEN, perhaps this could become a routine approach in future.

Improvements enabled by HPC with the current approach

Removal of the equivalent stoichiometric cloud assumption

HPC could allow a direct one-to-one coupling between the spatially-varying clouds at the dispersion stage and the simulated explosion cases, so that each snapshot of each cloud from the dispersion stage is simulated explicitly at the explosion stage, making the equivalent stoichiometric cloud assumption redundant.

In the immediate future the additional workload involved with this approach would probably preclude the potential value-adding extensions to the probabilistic approach discussed previously. However, it is recommended that comparisons are undertaken using HPC for a range of projects to investigate whether the outcome of a probabilistic assessment is sensitive to the equivalent stoichiometric cloud assumption. If it is demonstrated that this assumption is indeed fit for purpose then it can be retained for use with confidence, without attracting criticism in future.

Capturing deflagration detonation transition (DDT)

The explosion-specific codes are currently verified for subsonic deflagrations. They are not verified for DDT (deflagration-detonation transition) where the explosion becomes so intense that it becomes supersonic and, importantly, self-sustaining. Whilst DDT has traditionally been considered a rare event, there is a growing body of evidence, particularly following the Buncefield inquiry, to suggest that DDT may be more common than previously thought [12].

Understanding the onset of DDT is an active research topic and the use of HPC to investigate DDT numerically from first principles using general-purpose CFD codes may provide useful understanding of this phenomenon. Methods to capture the onset of DDT and the dynamics of the subsequent detonation might then be incorporated into the explosion-specific codes.

Beware of HPC!

Beware of modelling the wrong scenario extremely accurately

The statistician George Box said: *all models are wrong but some are useful* [13]. Over the past couple of decades there has been a continued effort to validate the explosion-specific CFD codes and whilst they could generally be improved, they are certainly useful and generally considered to be fit-for-purpose tools for undertaking probabilistic explosion assessments.

With modern HPC there is the prospect to more accurately model the underlying physics of the flow behaviour at each stage of the probabilistic analysis, but one has to ask – *is it really worth it?* Whilst it may be possible to simulate a high-pressure release more accurately with an adaptive, unstructured general-purpose code, or model an explosion from first principles by explicitly capturing congestion in an unstructured CFD mesh and simulating the detailed combustion physics at the flame front, would this actually add any additional value?

Given that there are already many assumptions and limitations required to make the probabilistic framework a manageable approach, and this is likely to remain the case even with HPC, and there will continue to be uncertainties relating to the density and distribution of small-scale obstructions in congested spaces, we must ask ourselves: *how accurately should we simulate what is potentially the wrong case?*

An HPC computing resource may allow a first principles approach to be pursued for a handful of explosion cases, but this approach is unlikely to be used within a probabilistic framework in the near future, where hundreds of individual simulated explosion cases may be required. Even if the most powerful super-computers were available and this approach was possible, it's probably not going to provide any more value to a project than using a simpler verified explosion-specific CFD code. However, this approach could provide additional understanding about the underlying physical mechanisms and DDT which could be used to improve the LPC-focussed distributed porosity methods currently employed by the explosion-specific CFD codes (which could also be incorporated into the general-purpose codes if their vendors were interested in doing so). But the LPC simplifications should be retained, rather than jettison this knowledge in favour of a significantly more computationally intensive first principles approach that is reliant upon ever more powerful HPC resources.

Beware of HPC induced complacency

Engineers should not be complacent and rely on HPC to forgo sound engineering judgement. In particular, the meshing tools now available allow users to surface wrap complex geometries and automate the meshing procedure with a few clicks of a mouse button. In our recent experience, sometimes these tools are used without any consideration as to whether a more efficient mesh could be constructed easily using more traditional mesh-building tools.

This was recently demonstrated when a general-purpose CFD code was used for a dispersion study for an offshore platform where, instead of using a porosity approach such as the ACE method which would have required a CFD mesh of a few million cells, the mesh attempted to capture all of the topsides congestion explicitly, resulting in an excessive mesh comprising around 60 million cells. Whilst there were impressive HPC resources available, which perhaps influenced this decision, the delivery of the study within the project timescale was put at risk and, more seriously, because the cells had been consumed in regions which were not of interest, it was not possible to refine the CFD mesh adequately in the regions that were of interest, even with the HPC resource. This was clearly apparent in the CFD predictions presented, as shown in Figure 7.

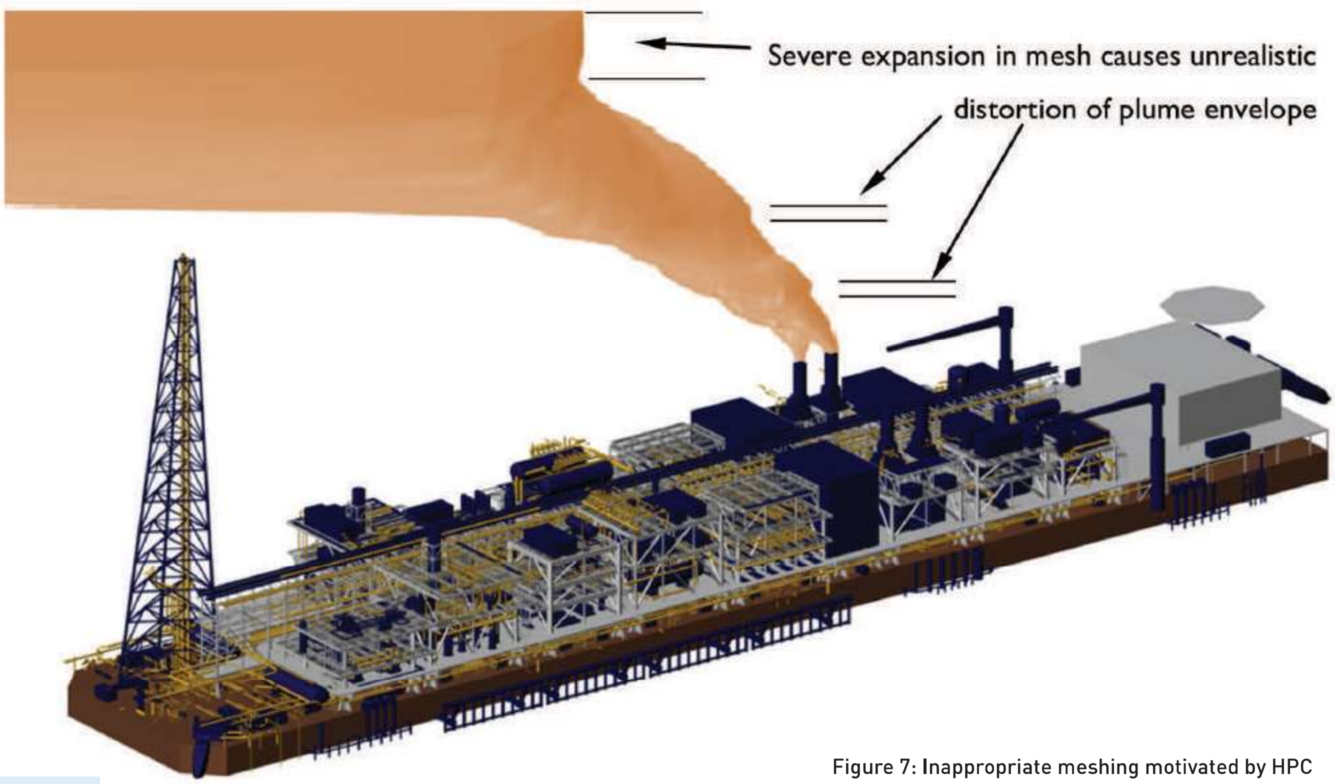

Figure 7: Inappropriate meshing motivated by HPC complacency prevented the CFD mesh from being adequately refined in the regions of interest

In Conclusion

HPC is a powerful enabling technology that can be exploited to successfully deliver probabilistic analyses. With this approach, a large number (hundreds or thousands) of cases can be simulated to allow the sensitivity of the CFD/FEA predictions to be understood with respect to a wide range of input parameters including the assumed boundary conditions and material properties. For the probabilistic explosion approach considered within this article the use of HPC can not only significantly reduce the time required to complete a study, it opens up new opportunities to improve the method which could add significant value in future.

Whilst this article has focussed upon the probabilistic framework for CFD-based explosion modelling, there are many other applications that could benefit from the probabilistic approach. Indeed, the reader is directed to another NAFEMS publication which describes a similar framework for FEA modelling [14] where HPC offers similar opportunities.

The successful exploitation of HPC for probabilistic studies relies on retaining a fit-for-purpose LPC approach for the individual simulations within the underlying dataset so that they each run quickly on the HPC host, together with the use of suitable SDM tools for automating the pre-processing and post-processing workflows. Retaining a fit-for-purpose LPC approach is critical – just because HPC enables the use of more complex methods and meshes it does not mean they are appropriate, and it is urged that engineers should not simply adopt an unnecessarily complex approach without first considering what is fit for purpose.

References

[1] Mercx WPM, *Modelling and experimental research into gas explosions, overall final report on the MERGE project*, Commission of the European Communities Report, Contract STEP-CT-011 (SSMA), 1993.

[2] Selby CA and Burgan BA, *Blast and fire engineering for topside structures - phase 2 (Final Summary Report)*, SCI Publication 253, The Steel Construction Institute, Ascot, U.K, 1998.

[3] Ely A, *The ACE method*, MSc at University of Leeds, 2004 (which can be download from www.abercus.com/SoftwareSolutions_ACEMethod).

[4 *Validation of a simplified approach to modelling supersonic gas release*, Howat G, Patel N and Roberts S, NAFEMS European conference on multi-physics simulation, UK, October 2014.

[5] Annex F of *Risk and emergency preparedness assessment*, NORSOK Standard Z 013, 2010.

[6] *IP Research Report: Ignition probability review, model development and look-up correlations*, Energy Institute, London, 2006.

[7] *Equivalent cloud methods used for explosion risk and consequence studies*, Hansen OR, Gavelli F, Davis SG and Middha P, Journal of loss prevention in the process industries 26, 511 527, 2013

[8] *Free Trial of High Performance Computing Services*, news article on Gexcon website (26 November 2015), www.gexcon.com/news/article/free-trial-of-high-performance-computing-services

[9] *FLACS HPC service: explosion modelling for industry*, news article on EPCC website (10 December 2015), www.epcc.ed.ac.uk/news/flacs-high-performance-computing-service

[10] *Enabling improved analysis data for blast protection of offshore structures*, Howell S, IMechE Offshore Engineering conference, Aberdeen, December 2015 (which can be download from www.abercus.com/News_20151207.aspx).

[11] *Risk-based structural response against explosion blast loads: systematic one-to-one CFD (FLACS) / NLFEA (Impetus AFEA solver) coupling to derive quantified response exceedence*, Salaun N, Hanssen AG and Nilsen PE, Chemical Engineering Transactions of the Italian Association of Chemical Engineers 48, 2016.

[12] *The next major hazard – controlling the onset of detonation*, Tam V and Johnson M, IMechE Offshore Engineering conference, Aberdeen, December 2015.

[13] *Robustness in the strategy of scientific model building*, Box GEP, in *Robustness in Statistics*, Launer RL and Wilkinson GN, Academic Press, pp 201–236, 1979.

[14] *Why do Probabilistic Finite Element Analysis?* Thacker BH, NAFEMS, 2008.

Appendix A - The Ace Method

The ACE method – a distributed porosity concept designed for general/unstructured CFD

The ACE method was developed to provide a consistent approach for describing the resistance to flow due to the small-scale obstructions which are abundant across the topsides of any offshore platform. The method is demonstrated below for a simple pipe bundle.

On the left hand side, the geometry of the pipe bundle is explicitly captured within the CFD mesh and the corresponding predicted pressure field is shown for cross flow (from left to right). On the right hand side a coarse mesh is used and the effect of the sub grid congestion is captured by the ACE method where equivalent resistance source terms are allocated in the momentum equations. A comparison of the two pressure predictions shows good agreement for the macro pressure behaviour but the mesh required for the ACE method, and the associated computational effort, is significantly reduced.

The ACE method is designed for any class of CFD mesh, including general unstructured meshes. The method was first implemented using the UDF functionality of the FLUENT CFD code (by Ansys). This implementation was completed by Alice Ely in 2004 during the course of her MSc at the University of Leeds. Her thesis is available for download from

www.abercus.com/SoftwareSolutions_ACEMethod

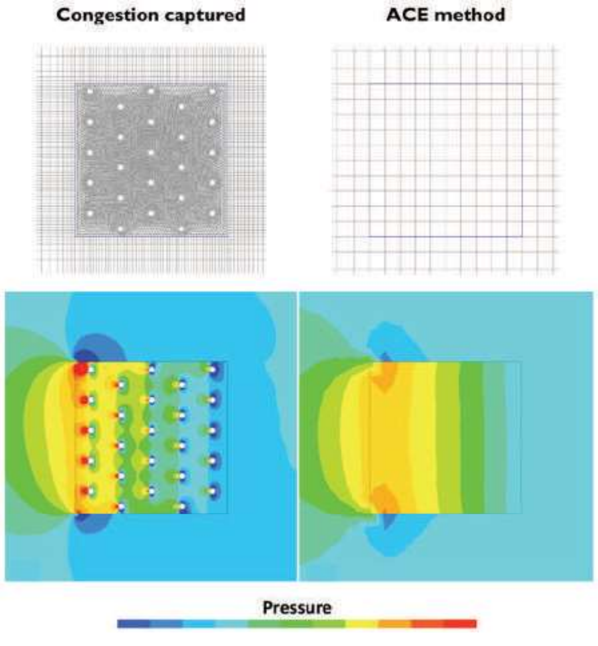

Appendix B - Simplification Of The Near-Jet Region

Simplification of the near-jet region for a high pressure compressible release

For high-pressure compressible releases, the near field region of the emerging jet has a high momentum and is typically not significantly affected by the background ventilation flow. Instead of explicitly capturing this region within the CFD model, it is usually advantageous to separate the near field and use another method to predict the flow, either using a classical calculation or a separate axisymmetric CFD approach. The emerging flow can then be introduced into the CFD model further downstream, at the end of the near field region (as shown below).

In order to capture the near field region using CFD a fine mesh discretisation may be required to properly capture the sequence of shocks as the discharged material expands to the ambient conditions, and the physics of the compressible flow need to be captured. By separating the near field region from the CFD model, the requirements for mesh discretisation are relaxed and compressibility effects can be neglected. It has been demonstrated that this approach, where the near-field is considered separately and only the far-field is simulated, will not significantly affect the predicted dispersion behaviour in the far-field region [4].

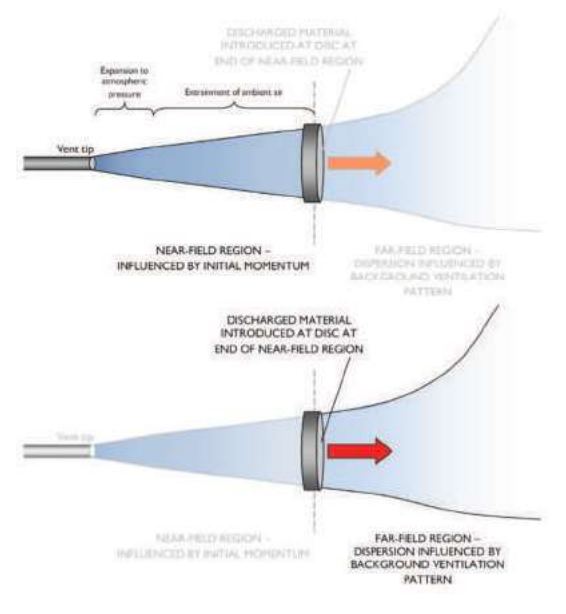

Appendix C -Equivalent Stoichiometric Volume

The equivalent stoichiometric volume

A hydrocarbon gas will only combust if its concentration in air falls within a certain band, between the lower flammable limit (LFL) and the upper flammable limit (UFL) – if the concentration of hydrocarbon gas is lower than the LFL there is insufficient fuel and the mixture is described as fuel lean, and if the concentration is above the UFL there is insufficient oxygen and the mixture is described as fuel rich. The LFL and UFL represent the concentrations where combustion can just be sustained, although the rate of combustion will be relatively low. At some point between the LFL and UFL will lie the stoichiometric concentration where the mixture of hydrocarbon gas and oxygen is just right for complete combustion. At around the stoichiometric concentration the rate of combustion and the ferocity of an explosion will be at its maximum.

For methane in air, the LFL, stoichiometric concentration and UFL are typically assumed to be 5%, 10% and 15% by volume respectively.

Appendix D - 3D Risk Assessment

Blast wall of interest, represented by a discretised array of monitor panels within the FLACS model

For this example, the 10-4/yr peak overpressure for the blast wall from the exceedence curve is 2 barg. Traditionally only this load would have been reported to the structural engineer and it would be assumed to apply uniformly over the entire blast wall. By considering the spatial variation, as shown in the contour plot to the right (which is an automated output from EXCGEN), it is apparent that the 2 barg 10 4/yr peak overpressure is the worst case load and that it is local to the bottom right side of the blast wall, with significantly lower 10 4/yr loading across the rest of the wall.

The contour plot for the 10 4/yr overpressure is assembled by constructing exceedence curves separately for each of the monitor panels within the blast wall, reading off the 10 4/yr overpressure at each monitor panel and then plotting this information in the form of the contour plot shown. The method can be applied to any risk acceptance/frequency cut off criteria and to any of the explosion loads of interest (overpressure, underpressure, positive/negative impulse, drag). EXCGEN simply automates this process.

The assumption for the basis of the structural design load has a significant impact upon the associated structural response of the blast wall. For this example, if the 2 barg peak overpressure is applied uniformly across the wall (which has been the traditional approach) the damage sustained by the wall is significantly more than the damage sustained by considering a pseudo-event using the spatially varying 10-4/yr overpressure shown in the contour plot [10].

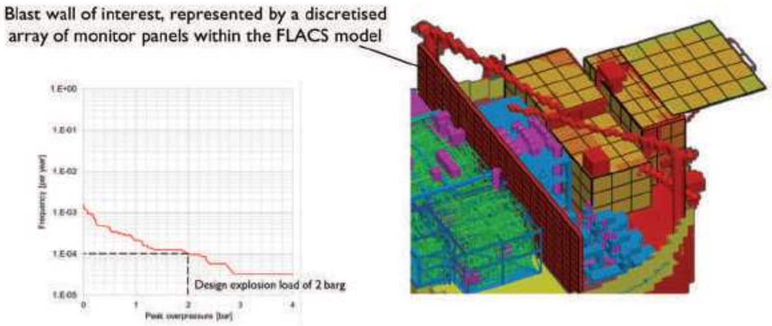

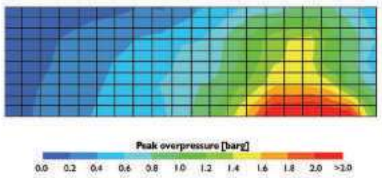
Contour plot of 10^{-4}/yr peak overpressure

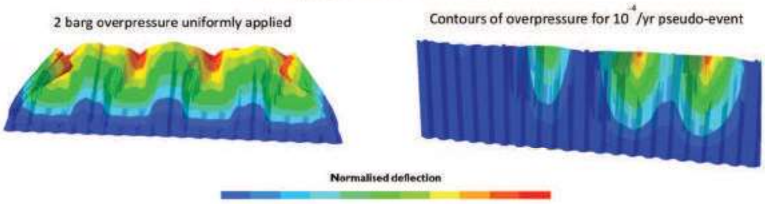
Contour plot of normalised deflection